REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
28-02-2011

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
01/03/2007 - 30/11/2010

4. TITLE AND SUBTITLE
Temporal Planning for Automatic Service Composition

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-07-1-0262

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Edmund H. Durfee

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The Regents of The University of Michigan
Div of Res Development & Administration
1058 Wolverine Tower
3000 South State Street
Ann Arbor, Michigan 48109

8. PERFORMING ORGANIZATION REPORT NUMBER
F017256 - 054979

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
David R. Luginbuhl
Air Force Office of Scientific Research
875 North Randolph Street
Suite 325, Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-OSR-VA-TR-2012-0182

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT : This project created, implemented, and evaluated distributed computational mechanisms for automating composition and scheduling of computer network-based services in applications where timing is important. For example, a commander might need intelligence to be gathered, processed, summarized, and visualized, but with timing preferences, such as that the results are available by a deadline but should be as up-to-date as possible. Preferences and constraints on timing impact which service providers are selected (depending on availability), which services they will provide (depending on the timing characteristics for different levels of service provision), and when each should begin and complete delivery of its promised service(s). Decision making is complicated by factors such as uncertainty over how long some services might take, competing service requests for scarce services, and inherent distribution of authority and private knowledge across service providers and requesters.
The project's technical accomplishments included novel extensions to temporal reasoning algorithms that allow computational agents to cooperatively uncover and reconcile timing constraints between their schedules quickly and with minimal loss of privacy. In addition, the researchers devised new distributed techniques for allocating tasks to service-providing agents that scale to larger service networks and avoid trying to formulate schedules in impossible (over-subscribed cases). Finally, the project defined a novel abstraction for coordinating services with uncertain durations, and has developed algorithms that have been shown to solve problems orders of magnitude faster than the prior state of the art, and to solve larger problems than has previously been possible.

15. SUBJECT TERMS
Automated Planning and Scheduling; Multi-agent Systems, Service-Oriented Computing

16. SECURITY CLASSIFICATION OF: U			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** Kathryn DeWitt
a. REPORT UU	**b. ABSTRACT** UU	**c. THIS PAGE** UU			**19b. TELEPHONE NUMBER** *(include area code)* 734-763-6438

Temporal Planning for Automatic Service Composition

Edmund H. Durfee, University of Michigan

1. Project Objectives:

This project's objectives were to create, implement, and evaluate distributed computational mechanisms for automating composition and scheduling of network-based services in applications where timing is important. For example, a commander might need intelligence to be gathered, processed, summarized, and visualized, but with timing preferences, such as that the results are available by a deadline but should be as recent as possible rather than having gone stale. Preferences and constraints on timing impact which service providers are selected (depending on availability), which services they will provide (depending on the timing characteristics for different levels of service provision), and when each should begin and complete delivery of its promised service(s). Further complications arise due to factors such as uncertainty over how long some services might take, competing service requests for scarce services, and inherent distribution of authority and private knowledge across service providers and requesters. The project's technical thrusts were to develop novel extensions to temporal reasoning, multiagent sequential decision-making, and distributed constraint optimization, which contribute collectively to solving such problems.

2. Summary of Significant Work Accomplished

The project made substantial contributions to the science and engineering of computational techniques for multiagent sequential decisionmaking, distributed constraint optimization, and temporal planning. Each of these is summarized below, where full descriptions of each of these contributions are available in the references cited.

2.1 Multiagent Sequential Decision-Making for Service Composition and Coordination

Agents that are cooperatively providing services to each other generally act in uncertain domains, where how long it will take to provide a service, and the quality of the result of a service provision, might not be fully predictable. Conventional methods to solve such problems draw on multiagent sequential decision-making techniques that explicitly coordinate the agents' joint policy decisions. These techniques are inherently susceptible to the curse of dimensionality, as the agents' state, action, and observation spaces grow exponentially with the number of agents. This project has made fundamental advancements to solving such problems by developing principled representations and algorithmic techniques that allow agents to coordinate at a more abstract level of *influences*. Intuitively, the idea is that an agent (such as an agent requesting a service) need not always know the full policy of an agent with whom it is cooperating (such as an agent providing a service), but instead needs to know only how the other's policy will materially influence its own plans, such as when the service will actually be provided, and not what other service provision requests will be fulfilled before and afterward [WD2008, WD2009a, WD2009b, WD2009c].

These ideas were developed in the recently-completed dissertation of Stefan Witwicki [W2011]. That work has derived a new complexity characterization of the joint policy coordination problem, combining several complementary aspects of weakly-coupled problem structure, including *agent scope size* (corresponding to the number of an agent's peers whose decisions influence the agent's decisions), *state factor domain size* (corresponding the space of (belief) states that an agent must model in order to plan optimal decisions), and *degree of influence* (corresponding to the proportion of unique influences that peers can feasibly exert). Studied separately, these aspects provide a language for describing various conditional independencies that may exist among the plans of a group of agents. Together, these three aspects define a three-dimensional landscape that can be used to quantify the advantage gained through exploiting a problem's interaction structure and, ultimately, to predict the amount of computation needed to solve the problem [WD2011]. For agents that model their world using a Decentralized POMDP (Dec-POMDP), a bound on the worst-case computational complexity of optimal planning can be derived as follows:

$$O\left(\underbrace{\mathrm{EXP}(\overbrace{\mathbb{X}_i^{max}}^{\text{state factor scope size}})}_{\text{complexity of best response}} \cdot \underbrace{n \cdot \overbrace{(\overbrace{d_{\mathbb{P}}}^{\text{degree of influence}}\|\Pi_i^{max}\|)}^{\text{number of unique influences}}{}^{\overset{\approx\ \text{agent scope size - 1}}{\omega}}}_{\text{number of best responses}} + \underbrace{C_{\mathbb{P}} \cdot n \cdot (d_{\mathbb{P}}\|\Pi_i^{max}\|)^{\omega-1}}_{\text{overhead of enumerating influences}}\right)$$

where n is the number of agents. Weakly-coupled agents will tend to have a smaller state factor scope size (fewer features of each others' states that they can affect), a smaller degree of influence (fewer different ways that their planned actions can affect each other), and a smaller agent scope size (fewer agents that can affect any particular agent). More importantly, however, capturing these aspects in a single formula helps explain and predict performance in situations where agents coupled to different degrees along the different dimensions.

To emphasize weakly-coupled structure, the research conducted in this project introduced a (transition-dependent decentralized POMDP) model that efficiently decomposes into local decision models with shared state features. For instance, features that are controlled by one agent, but that directly affect the consequences of another agent's actions, such as happens between agents requesting and providing a particular service application, are included in both agents' models. From the perspective of the affected agent, these are referred to as nonlocally-controlled features. In essence, the conventionally-specified Dec-POMDP model has been decoupled into a set of local POMDPs tied to one another by the transition-dependence among their nonlocally-controlled features; this new model is thus referred to as a Transition-Decoupled POMDP (TD-POMDP). In comparison with related models, the TD-POMDP achieves an effective balance in its articulation of exploitable structure and its loss of generality [WD2010a, WD2010b, WD2010c, W2011].

With the TD-POMDP model structure thus defined, interagent influence may be characterized quite simply as the expected transition probabilities of nonlocally-controlled features. Since these probabilities are the only components of an agent's local model that may vary with the behavior of its peers, entire peer policies can be abstractly summarized by the influences they entail, and the corresponding probabilities incorporated into the transition model of a single-agent POMDP that serves as the agent's local decision model. The transition probabilities associated with a particular influence can be encoded with a probability distribution $Pr\ (n_j|f_1;f_2,\ldots)$, where n_j are new values of nonlocally-controlled features conditioned on previous values of various state features f_1, f_2, etc. This work has proven that, for any TD-POMDP, the influences for the system of agents can be jointly specified with a Dynamic Bayesian Network (DBN) containing only (variables representing the past and present values of) shared state features. This project's investigators have also proven that this influence representation is sufficient for optimal coordination, through the use of an influence space search methodology.

In essence, the joint policy formulation problem has been decoupled into the subproblems of (1) proposing influences, (2) evaluating influences, and (3) computing optimal policies around influences. A mixed-integer linear programming (MILP) methodology has been developed for solving each of these subproblems. In contrast with prior approaches geared towards enforcing interacting behavior, this novel methodology enables an agent to determine whether a desired influence is feasible, if so to compute the optimal local policy that is constrained to exert the influence, and to completely avoid any tuning of parameters associated with influence enforcement [WD2010a, WD2010b, W2011].

The primary advantage of working in the influence space is that there are potentially significantly fewer feasible influences than there are policies. Blending prior work on decoupled joint policy search and constraint optimization, the investigators have developed influence-space search algorithms that, for problems with a low degree of influence, compute optimal solutions orders of magnitude faster than policy-space search. When agents' influences are constrained, influence-space search also outperforms other state-of-the-art optimal solution algorithms.

The graphs below are examples from a rigorous empirical comparison of the optimal influence-space search (OIS) algorithm against three other state-of-the-art optimal solution algorithms tailored for specialized Dec-POMDP problems. As demonstrated, when agents' windows of interaction are small, indicating that agents' influences are most heavily constrained, influence abstraction is able to exploit this structure to gain exponential advantage (on average) over each of the other algorithms, which are less effective at exploiting this form of weakly-coupled structure [WD2010a, WD2010b].

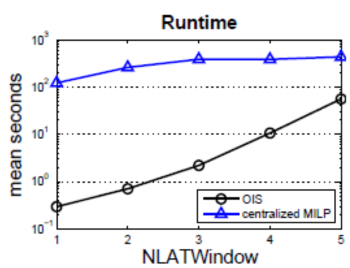

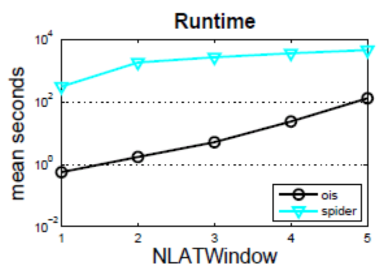

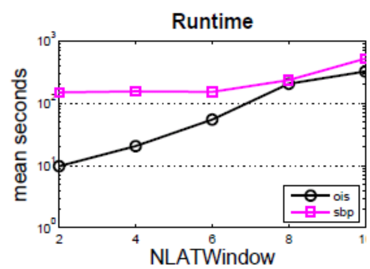

Moreover, by exploiting both degree of influence and agent scope size in a bucket-elimination variation [W2011] of the OIS algorithm (labeled be-ois in the graph below, compared to a depth-first OIS implementation), the investigators have demonstrated scalability, substantially beyond the reach of prior optimal methods, to teams of 50 weakly-coupled transition-dependent agents, as shown below.

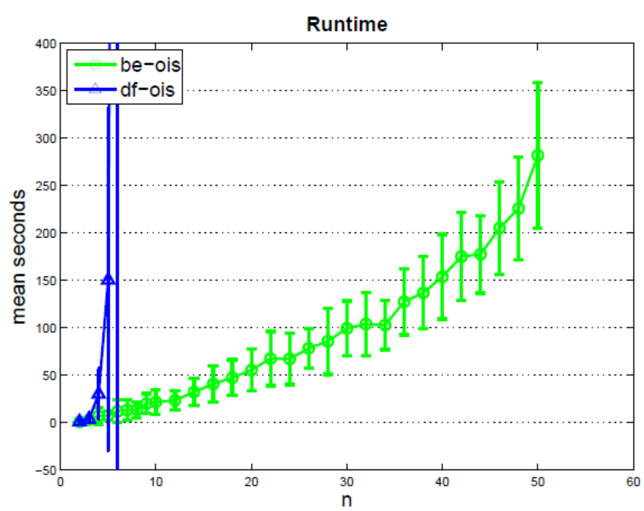

Extensions

During the course of this project, two Masters students have done research that has extended the ideas described above to increase the range of problems that can be tackled. One student, Anna Chen, has investigated how the concept of “influence”, and especially the notion that agents make commitments about how they will influence each other, extends to problems where the agents might be uncertain, at the outset of execution, about what states of the world are the most important to be reached. Such problems arise in planning and scheduling services, because in the midst of executing a planned set of services, new (high priority) service requests might arrive.

Her work [CDSW2011] has characterized the conditions under which agents can still maximize their expected joint performance by optimizing based on the mean-reward, rather than having to consider every possible reward function separately. It has also developed an iterative greedy algorithm for circumstances where mean-reward is not optimal, and has empirically demonstrated that the algorithm can run considerably faster than does reasoning about all possible reward functions, with only modest sacrifices in quality. Intuitively, the idea is that if a service provider "hedges" its commitments to retain some degree of local slack, then in expectation the joint performance can improve because it can be responsive to emergent opportunities. Beyond this intuition, though, is the question of "how much" slack is the right amount to introduce, and the work provides some preliminary answers to that question.

The investigators' work on improving tractability of finding joint policies for service provision has also identified some promising approximation techniques. They showed that approximately optimal commitments for service provision could be computed in a greedy manner as a form of distributed binary search [WD2008, WD2009a], or by constraining the number of time points considered by the search[WD2009b, WD2009c].

They have also, in the Master's project of Jason Sleight, analyzed the degree to which approximating durational uncertainty can be effective [S2011]. Specifically, that work has emphasized reducing the branching factor of possible future trajectories by reducing the number of different possible durations that a service might take down to a smaller number. That work has proven that the search for a duration approximation should never consider approximate durations that do not correspond to actual possible durations, and that inexpensive error metrics can fruitfully act as proxies for the expected loss of utility that such approximations will incur. It has also developed a polynomial time (dynamic programming) algorithm for quickly determining the appropriate approximation.

In addition, that work has demonstrated the importance of characterizing the execution-time behavior of an agent when determining the "best" approximation, in terms of what the agent does when the actual execution of a service takes a duration that was excluded from the approximation. The work has shown that for more passive execution-time responses, such as having the agent idle unless/until the natural evolution of the world leads back to an expected state, a simple parameterization of the dynamic programming algorithm can formulate an approximation that can better account for this execution-time behavior. The graph below (left) shows how different settings of the parameter (alpha) affect the loss of expected value as the number (K) of durations out of the set of actual discrete durations is increased from 1 to all 10, and shows the algorithms computation time as K grows (right). These types of performance-cost profiles provide information to a system developer to strike the right tradeoff for a particular application.

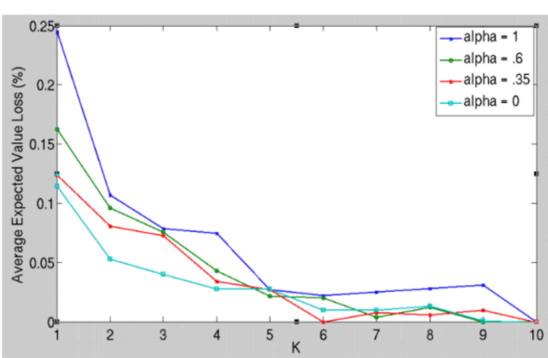

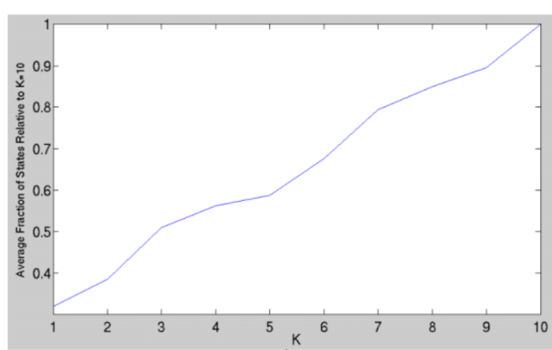

2.2 Distributed Constraint Optimization for Service Assignment

A danger in conducting temporal planning for service composition arises when the number of requests for a particular service could outstrip the capabilities of the provider of that service. Faced with such an “over-constrained” problem, a temporal planner can spend considerable time searching for a fully satisfying plan by trying different orderings and timings of service provisions, when in fact the problem is inherently unsolvable. To address this case, the Masters degree project of Christopher Portway developed a preprocessing step for efficiently detecting over-constrained situations, and for (heuristically) indentifying an approximately optimal subset of requests that could be collectively achieved with the provider’s limited resources.

The core idea behind the approach has been to cast the problem of deciding on service requests to fulfill as a distributed constraint optimization problem (DCOP). A variety of algorithms for solving DCOPs have been developed in the past; because the objective in this case is to optimize the value of the service requests satisfied within the constraints of the providers’ time constraints, the approach adopted in this project has built off of the multiply-constrained (MC-)DCOP framework developed by other researchers. However, that framework required significant augmentation to support modeling single agents each reasoning over multiple variables (requests), leading to the development of a new Multiple Variable (MV-)MC-DCOP [PD2009]. Further improvements to solution quality can be gained by including in the formulation the temporal (precedence) constraints that occur when services sequentially chain their contributions to yield an overall response to a user’s needs, leading to the Ordered (O-)MV-MC-DCOP [PD2010a, PD2010b].

These extended frameworks were implemented in the multiagent system as (O-)MV-MC-MGM, extending the efficient (but only approximately optimal) MC-MGM DCOP approach, that was developed by others, to now handle multiple variables and ordering constraints. Each of the service providers and requesters engage in the (O-)MV-MC-MGM protocol to converge on which subset of services to attempt to schedule, and only once this trimming to the problem has been done do these agents engage in the more costly search for joint influences and policies as described in Section 2.1. Empirically, the (O-)MV-MC-MGM preprocessing was applied to larger problems, and was shown to be able to improve expected reward and to reduce overall computation time compared to solving the problems using the approximate methods described in Section 2.1. An example of the results from [PD2010b] is shown below, which also includes the performance of a centralized mixed-integer linear program (MILP) technique for solving the MV-MC-DCOP.

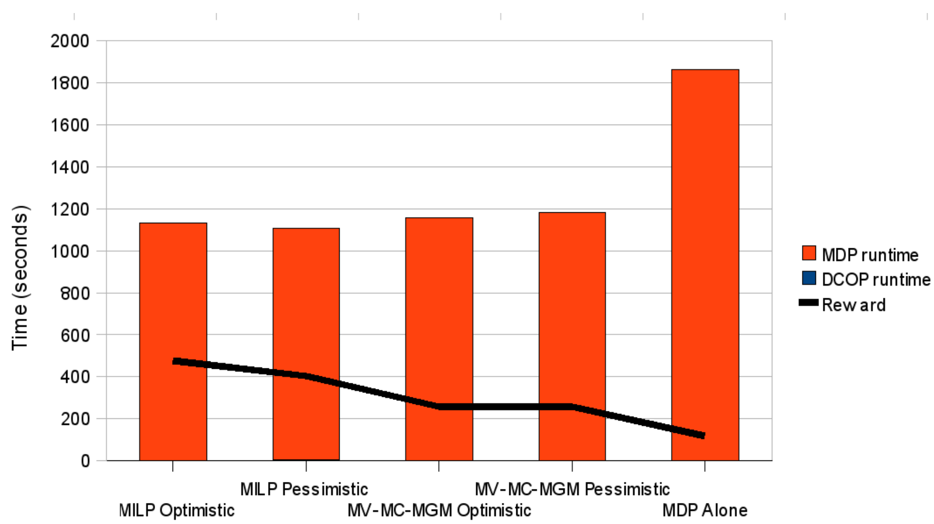

2.3 Improved Multiagent Selection and Scheduling Algorithms

Selection and scheduling of services (or activities more generally) are intertwined problems. That is, different selection choices can introduce different temporal constraints (e.g., "better" services might take longer to provide), and timing constraints can limit selection choices (e.g., an impending deadline rules out the "best" but longest duration service). For service composition problems involving services with more controllable/deterministic durations, the doctoral research of PhD student Jim Boerkoel has formulated the problem as a Hybrid Scheduling Problem (HSP). An HSP is comprised of a traditional finite-domain constraint satisfaction problem (CSP)—where service requests are variables, for example, and a service provider value needs to be selected for each—and a traditional disjunctive temporal problem (DTP)—where start and end times of particular service provisions must be assigned. These problems are coupled by "hybrid constraints" that express relationships between selection and scheduling assignments, such as that if service provider A is selected to satisfy a request, then the timing of when the request will be satisfied is constrained by provider A's prior commitments.

The research conducted in this project been developing techniques for automatically deriving and expressing critical implied constraints in HSPs that in turn enable state-of-the-art constraint propagation techniques to rapidly converge on satisfying solutions. The investigators have also developed the first partially and fully decentralized triangulation-based algorithms for solving multiagent temporal constraint problems, and demonstrated their efficacy. These two thrusts are summarized below.

Hybrid Constraint Tightening

Hybrid Constraint Tightening (HCT) is an algorithm for preprocessing an HSP formulation, applying constraint compilation principles to reformulate hybrid constraints by lifting information from the structure of an HSP instance [BD2008]. These reformulated constraints elucidate implied constraints between the CSP and DTP subproblems of an HSP earlier in the search process and can lead to significant search space pruning. Despite the computational costs associated with applying the HCT preprocessing algorithm, HCT leads to orders of magnitude speedup when used in conjunction with off-the-shelf, state-of-the-art solvers, as compared to solving the same problem instance without applying HCT. The investigators have conducted a systematic exploration of the properties of HSPs that influence HCT's efficacy, and have quantified empirically the conditions under HCT is particularly effective [BD2009].

Multiagent Algorithms for Temporal Reasoning and Decoupling

Activities such as satisfying service requests inherently link the schedules of different agents together by imposing *interagent* temporal constraints, such as when a service requester must wait to proceed on its next task until after the service provider has returned the requested result. Each agent also has *intra-agent* (internal) temporal constraints, among which is the constraint that if the agent is performing a task, then the start time of any other task it plans to do must be no sooner than the end time of the task it is performing. As part of this project, the investigators have formally defined the MaSTP – multiagent simple temporal problem – as a multiagent extension of the traditional STP representation that also accounts for constraints between agents' schedules, and hence the definition of what portions of each agent's STP is private to that agent, and which portions other agents STPs it must necessarily be aware of [BD2010a, BD2010b].

Prior research on solving problems that can be encoded in the MaSTP formulation have instead modeled the problem in a single, centralized way, and thus not only fail to scale well to larger multiagent systems, but also require full revelation of private scheduling information. Prior solution algorithms thus require centralizing the problem representation at some ``coordinator'' who calculates a (set of) solution schedule(s) for all, and such algorithms can incur unacceptable computational, communication, and privacy costs for problems like service composition where problems are inherently distributed and where unnecessarily propagating awareness of what services are making requests of what other services might be risky. That is, agents that specify problems in a distributed fashion might reasonably expect some degree of privacy, and demand greater latitude for rapid scheduling changes than can be supported in a centralized system.

This project has developed new, *distributed* algorithms for finding and maintaining a (set of) solution(s) for the MaSTP. The investigators have proven the correctness, privacy implications, and runtime properties of each of these algorithms. They have also empirically evaluated the algorithms' costs in terms of both time and communication.

The high-level approach is to decompose problems into *n* locally-independent subproblems that each of the *n* agents can solve concurrently and privately, and one shared subproblem for which the agents must work together to solve. This decomposition divides variables based on whether or not they are involved in external constraints and so exploits the natural, loosely-coupled problem structure of many real-world problems, such as service composition. This partitioning is described in detail for the MaSTP in [BD2010a], and is used to prove important privacy properties for the algorithms and to empirically demonstrate how this structure influences algorithm performance. Solving the MaSTP is an important precursor for solving the multiagent versions of more complicated scheduling formulations such as DTPs and HSPs, which often require quickly evaluating partial, candidate assignments in the form of *component STPs*.

The algorithm to solve the MaSTP employs a variable elimination procedure, where each agent privately and concurrently eliminates its private local variables first, and then coordinates with other agents to eliminate its externally constrained (shared) variables. The algorithm then performs a reverse pass that calculates the full set of possible joint solutions. This distributed algorithm demonstrates impressive levels of speedup over comparable centralized approaches, especially on weakly-coupled problems. The graph below compares the MaSTP algorithm where agents process their private problems concurrently, with a variation that processes the private problems sequentially. These are both compared to solving the problem in the standard centralized way. As the proportion of interagent to intra-agent constraints is varied from low (left) to high (right), the MaSTP algorithms display a speedup of between 5 and 2 orders of magnitude!

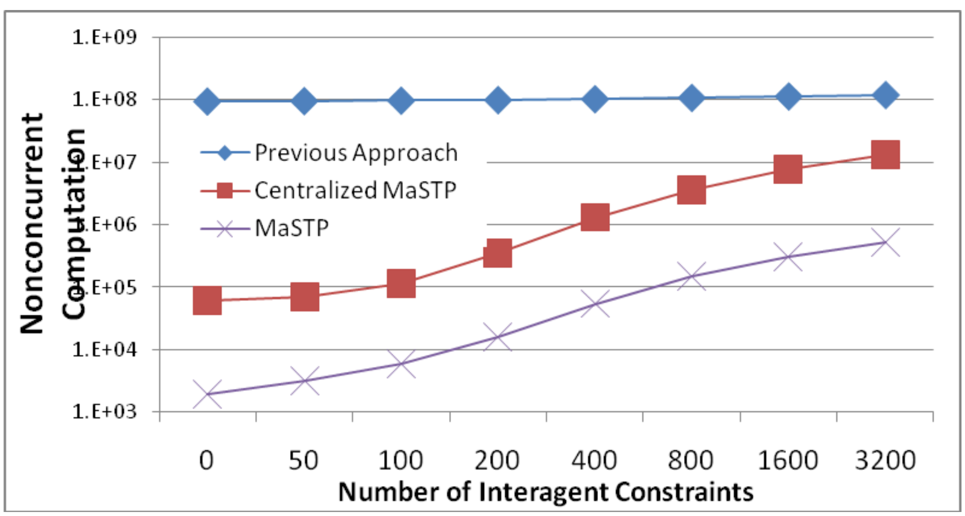

The resulting solution to the MaSTP provides agents with the maximal amount of flexibility permitted by their constrained problems, and can support the further refinement of their schedules to converge on coordinated activities. However, this further refinement requires that, each time an agent assigns one of its timepoint variables, all other agents must wait until that decision has been propagated across the network before any other agent can make a decision. Without waiting for such propagation to complete risks agents making incompatible assignments that render the joint scheduling problem unsolvable.

An alternative that has been developed in this project is to allow agents to heuristically *decouple* themselves from the subproblems of other agents. Informally, a decoupling (e.g., a temporal decoupling) is defined in terms of locally independent sets of solutions that, when combined, form a solution to the original multiagent constraint problem. Thus, the goal for each agent is to make search decisions (e.g., impose new intraagent constraints) that render interagent (external) constraints moot. For example, if a service requester and provider agree on a time at which the service provision must finish, then they can schedule incorporate this more precise constraint into their local models and know implicitly that their interagent constraint will be satisfied.

The investigators have augmented the MaSTP algorithm with this ability to heuristically decouple agents' problems, resulting in a solution to the multiagent temporal decoupling problem (MaTDP) [BD2011]. Specifically, during the reverse phase of MaSTP execution, agents introduce new constraints into their shared STP that temporally decouple their local subproblems. Once all subproblems are decoupled, the reverse phase continues into the agents' private problems. Finally, the investigators have also developed an algorithm that makes a forward pass through the shared problem after the decoupling to relax unnecessarily tight constraints before completing the reverse pass through the agents' private problems.

In the graph below, the computational time for solving the MaTDP plus relaxation is shown to be several orders of magnitude faster than the previous best algorithm for solving the TDP (labeled TDP) which did so in a centralized way. A centralized (sequential rather than concurrent) version of the MaTDP algorithm is also tested, as is the MaSTP algorithm.

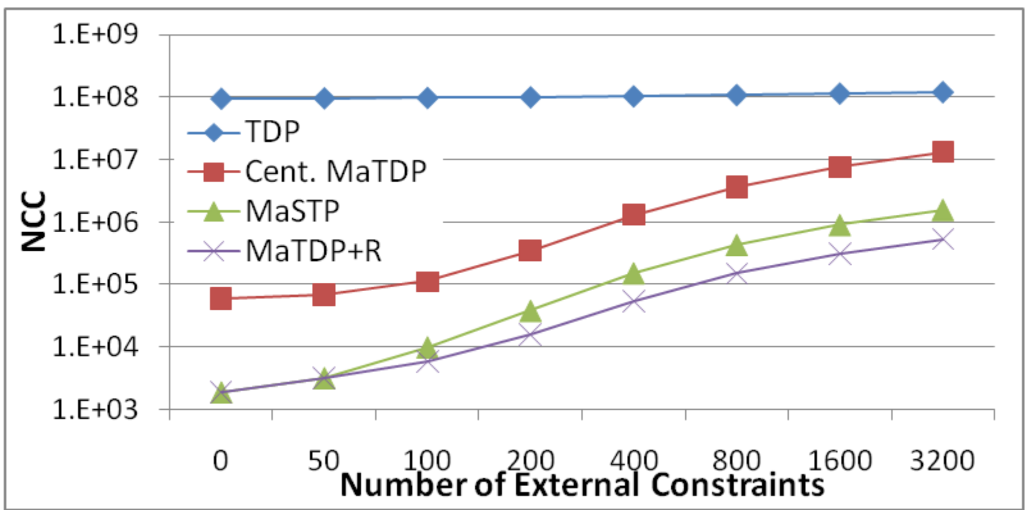

Further, using the traditional measure of performance for decoupling, which measures the rigidity in a decoupled schedule (where finding a decoupling that minimizes rigidity is good), the MaTDP+R technique performs nearly as well as the state-of-the-art centralized TDP algorithm, despite the MaTDP+R heuristic being much simpler and being computed in a distributed manner. In the table below, the minimal rigidity for the non-decoupled case is compared to rigidity for TDP and MaTDP+R at levels of external constraints (from the graph above) of 50, 200, and 800.

Algorithm	N=50	N=200	N=800
No Decoupling	0.418	0.549	0.729
TDP	0.482	0.668	0.865
MaTDP+R	0.496	0.699	0.886

3. Summary

This project developed and evaluated theoretically sound and practically relevant techniques for multiagent sequential decision-making, distributed constraint reasoning, and multiagent scheduling and planning. These results are applicable to solving problems of temporal planning for service composition, and have broader promise for other applications, such as for finding and scheduling human experts for collaboration to solve complex problems [DBS11].

4. Project Personnel

Edmund H Durfee, Professor (PI)
Stefan Witwicki, graduate student (PhD completed January 2011)
James Boerkoel, graduate student (PhD expected Summer 2012)
Christopher Portway, graduate student (Masters degree completed May 2010)
Jason Sleight, graduate student (Masters degree December 2010, current PhD student)
Inn-Tung (Anna) Chen, graduate student (Masters degree expected May 2011)

5. References (Project Publications)

With the exception of paper [DBS11], which was supported as part of an SBIR working with Intelligent Automation Incorporated, all of the referenced papers below were supported in part by this grant.

[BD2008] James C. Boerkoel Jr. and Edmund H. Durfee. "Hybrid Constraint Tightening for Solving Hybrid Scheduling Problems." In *Proceedings of 23rd AAAI Conference on Artificial Intelligence (AAAI-08)*, pages 1446-1449, July 2008.

[BD2009] James C. Boerkoel Jr. and Edmund H. Durfee. "Evaluating Hybrid Constraint Tightening for Scheduling Agents." In *Proceedings of Eighth International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-09),* pages 673-680, May 2009.

[BD2010a] James C. Boerkoel Jr. and Edmund H. Durfee. "A Comparison of Algorithms for Solving the Multiagent Simple Temporal Problem." In *Proceedings of the 2010 International Conference on Automated Planning and Scheduling (ICAPS-10)*, pages 26-33, May 2010.

[BD2010b] James C. Boerkoel Jr. and Edmund H. Durfee. "Partitioning the Multiagent Simple Temporal Problem For Concurrency and Privacy (Extended Abstract)." In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems* (*AAMAS-2010*), Toronto, pages 1421-2, May 2010.

[BD2011] James C. Boerkoel Jr. and Edmund H. Durfee. "Distributed Algorithms for Solving the Multiagent Temporal Decoupling Problem." To appear in *Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-11)*, May 2011.

[CD2009] Jeffrey S. Cox and Edmund H. Durfee. "Efficient and Distributable Methods for Solving the Multiagent Plan Coordination Problem," *Multiagent and Grid Systems* 5(4): 373-408, 2009. (Special Issue on Coordinating Agents' Plans and Schedules.)

[CDSW2011] Inn-Tung Anna Chen, Edmund Durfee, Satinder Singh, and Stefan Witwicki. "Influence-Based Multiagent Planning under Reward Uncertainty." To appear in Working Notes of the AAMAS-11 Workshop on Multiagent Sequential Decision Making, May 2011.

[D2008] Edmund H. Durfee. "Planning for Coordination and Coordination for Planning." (Invited Abstract) 2008 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology, pages 1-3, December 2008.

[DBS2011] Edmund H. Durfee, James C. Boerkoel, Jr., and Jason Sleight. "Comparing Techniques for the Semi-Autonomous Formation of Expert Teams." To appear in the *Proceedings of the 2011 International Symposium on Collaborative Technologies and Systems (CTS 2011).*

[PD2009] Christopher Portway and Edmund H. Durfee. "The Multi Variable Multi Constrained Distributed Constraint Optimization Framework." *Working Notes of the IJCAI 2009 Workshop on Distributed Constraint Reasoning,* pages 102-116, July 2009.

[PD2010a] Christopher Portway and Edmund H. Durfee. "The Multi Variable Multi Constrained Distributed Constraint Optimization Framework (Extended Abstract)." In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems* (*AAMAS-2010*), Toronto, pages 1385-6, May 2010.

[PD2010b] Christopher Portway and Edmund H. Durfee. "Ordered Multi-Variable Multi-Constrained Distributed Constraint Optimization Framework." In *Proceedings of the 2010 Intelligent Agent Technology Conference (IAT10)*, September 2010.

[S2011] Jason Sleight. "Approximating Stochastic Task Durations for State Space Reduction in an MDP." Masters' project. University of Michigan.

[W2011] Stefan J. Witwicki. *Abstracting Influences for Efficient Multiagent Coordination Under Uncertainty.* Doctoral Dissertation, University of Michigan, January 2011. Available at http://www.umich.edu/~witwicki/bib/dissertation/.

[WD2008] Stefan J. Witwicki and Edmund H. Durfee. "Commitment-Based Service Coordination." In *Proceedings of the International Workshop on Service-Oriented Computing: Agents, Semantics, and Engineering (SOCASE-08)*, Springer Lecture Notes in Computer Science, pages 134-148, 2008.

[WD2009a] Stefan J. Witwicki and Edmund H. Durfee. "Commitment-Based Service Coordination," *International Journal on Agent Oriented Software Engineering,* Special Issue on Service-Oriented Computing: Agents, Semantics, and Engineering. 3(1):59-87, 2009.

[WD2009b] Stefan J. Witwicki and Edmund H. Durfee. "Flexible Approximation of Structured Interactions in Decentralized Markov Decision Processes (Extended Abstract)." In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems* (*AAMAS-2009*), pages 1251-2, Budapest, Hungary. May 2009.

[WD2009c] Stefan J. Witwicki and Edmund H. Durfee. "Flexible Approximation of Structured Interactions in Decentralized Markov Decision Processes." *Working notes of the AAMAS-09 Workshop on Multiagent Sequential Decision Making*, May 2009.

[WD2010a] Stefan J. Witwicki and Edmund H. Durfee. "Influence-Based Policy Abstraction for Weakly-Coupled Dec-POMDPs." In *Proceedings of the 2010 International Conference on Automated Planning and Scheduling (ICAPS-10)*, pages 185-192, May 2010.

[WD2010b] Stefan J. Witwicki and Edmund H. Durfee. "From Policies to Influences: A Framework for Nonlocal Abstraction in Transition-Dependent Dec-POMDP Agents (Extended Abstract)." In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems* (*AAMAS-2010*), Toronto, pages 1397-8, May 2010.

[WD2010c] Stefan J. Witwicki and Edmund H. Durfee. "Influence-Based Policy Abstraction for Weakly-Coupled Dec-POMDPs." In *Working Notes of the AAMAS-10 Workshop of Multiagent Sequential Decision Making*, May 2010. (Extended version of paper that appeared in ICAPS10.)

[WD2011] Stefan J. Witwicki and Edmund H. Durfee**. "**Towards a Unifying Characterization for Quantifying Weak Coupling in Dec-POMDPs." To appear in *Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-11)*, May 2011.

[WuD2010] Jianhui Wu and Edmund H. Durfee. "Resource-Driven Mission-Phasing Techniques for Constrained Agents in Stochastic Environments," *Journal of Artificial Intelligence Research (JAIR)*, volume 38, pages 415-473, July 2010.